Intelligent Filtering for Augmented Reality

Sabrina Sestito*, Simon Julier, Marco Lanzagorta and Larry Rosenblum

Advanced Information Technology (Code 5580)
Naval Research Laboratory
Washington DC
(* on attachment from the Aeronautical and Maritime Research Laboratory, Defence Science and Technology Organisation, Melbourne, Australia)

KEYWORDS: Augmented Reality, Intelligent Systems, Databases

ABSTRACT:
Recent developments in computing hardware have begun to make mobile and wearable Augmented Reality (AR) systems a reality. With this new freedom, AR systems can now be used in a very wide range of applications including disaster relief, localization and repair of utilities and even as an assistant for tourists walking through unfamiliar historical sites. This paper considers the use of AR to assist with the task of warfighting in an urban environment. Urban environments are compact, complicated, and can be highly dynamic. Any successful AR system must overcome a host of challenges including the need for a robust tracking systems and wearable hardware and must present the information to the user an intuitive and informative manner. This paper considers the problem of designing a user interface which avoids information overload through automatically managing the grahpical context which is displayed to the user. We describe the paradigm of an information filter - a decision mechanism that uses the user's location, the user's current goal and the properties of objects within the environment to deduce what information should be displayed.

1. Introduction

It is expected that many future military operations will occur in urban environments [CFMOUT-97]. These present many unique and challenging conditions for the warfighter. The environment is extremely complicated and inherently three-dimensional. Above street level, buildings serve many purposes (such as hospitals or communication stations) and can harbor many risks (such as snipers or mines) which can be located on many floors. Below street level, there can be a complex network of sewers and tunnels. The environment can be very cluttered. Narrow streets restrict line of sight and make it difficult to plan and coordinate group activities. The environment can be highly dynamic and in constant flux. Threats (such as snipers) can continuously move and the structure of the environment itself can change. For example, a damaged building can fill a street with rubble, making a once safe route impassable. These difficulties are compounded by the need to minimize the number of civilian casualties and the amount of damage to civilian targets.

These and other difficulties have lead the Concepts Division of the Marine Corps Combat Development Command to conclude that

"Units moving in or between zones must be able to navigate effectively, and to coordinate their activities with units in other zones, as well as with units moving outside the city. This navigation and coordination capability must be resident at the very-small-unit level, perhaps even with the individual Marine" [CFMOUT-97]. Therefore, the success of a military operation in an urbanized environment depends crucially on being able to provide navigation and coordination information at the individual marine level.

A number of research programs have explored the means by which navigation and coordination information can be delivered to the dismounted soldier. Many of these approachs are based on handheld maps (e.g., an Apple Newton), or opaque head mounted displays (HMD). For example, the Land Warrior program introduced a headmounted display which combined a map and a “rolling compass” [Gumm-98]. Unfortunately, these methods have a number of limitations. They obscure the user’s field of view and do not truly represent the three dimensional nature of the environment. To overcome these problems, we propose the use of a *mobile augmented reality* system.

A mobile augmented reality system consists of a computer, a tracking system and a see-through head mounted display. The system tracks the position and orientation of the user and superimposes, within the user’s field of view, graphics and annotations which are aligned with objects in the environment. This approach has many advantages. Information can be presented in an intuitive manner and integrated directly with the environment. For example, the name of a building would appear as a “virtual sign post” attached directly to the side of the building. To explore the feasibility of such a system, the Naval Research Laboratory (NRL) is developing a prototype augmented reality system known as BARS, the Battlefield Augmented Reality System. This system will network multiple, outdoor, mobile users together with a command centre. To achieve this goal many challenges must be overcome [Julier-99]. These include hardware issues (accurate and robust tracking, high performance head mounted displays and compact wearable computers), software issues (distributed, shared three-dimensional environments) and human computer interaction (what information should be displayed and how).

This paper examines the proble computer interaction – how does the system determine, from the user’s context, what graphical information should be displayed? This is extremely important because information overload is a significant potential problem.As a user moves through the environment, their context can change dramatically depending on their position and current intent. The amount of information that can be shown to a user in a virtual world can be overwhelming. To alleviate this problem, the mobile system must sort and prioritize information so that only the features which are “most relevant” to the user’s current state and location should be shown. In this paper, we argue that this is best achieved through an 'intelligent' filter which determines what information is relevant to the user at a particular time. However, the design of such a system cannot be addressed separately from the model of the environment which is maintained by the system.

The structure of this paper is as follows. Section 2 summarises the database structure used in the Battlefield Augmented Reality System (BARS). The information filter is described in Section 3. The BARS prototype is briefly described in Section 4 and future work described in Section 5. Conclusions are drawn in Section 6.

2. The BARS Database

2.1 Object Model of the Environment

The focus of BARS is to deliver information to a dismounted warfighter operating in an urban environment. The system has to potentially provide information about very fine-grained features such as a particular door in a particular building. Therefore, the database was designed to be object-oriented. The environment is assumed to be populated by large numbers of objects which have certain logical and physical relationships between one another. All objects share the following properties:

- name
- position
- type
- *importance* vector. This vector is described in more detail in Section 3.3.2

The top of the hierarchy is the *City* which defines the region within which the user is operating. The City possesses a number of Forces (mobile objects which can be enemies or friends) and a set of “regions”. Each “region” corresponds to a feature such as a building or a street and has the property that it contains discrete features which are physically grouped together. Building objects, for example, possess walls, windows and doors.

2.2 Databases

The database must be sufficiently expressive that it can be used to render a scene with a high degree of accuracy, or to represent it as a set of individual components which can be shown separately. Furthermore, the distribution of the databases must be such that individual mobile clients will operate even if the connection with the base station is disrupted. To achieve these goals, the system actually uses three different types of databases which are linked to one another – the Object Database, the Visual Database and the User Database. The Object Database contains all of the information known about the urban environment in a symbolic form; for instance, the names of buildings and locations of snipers. The Visual Database contains the geometrical/visual data (textures and polygons) necessary to visualize the environment. The third database is the Users' Database; it contains the information that is displayed on the users' HMD (e.g., military icons or names of buildings). The current system, which consists of a single mobile user, merges the Object and User Database together.

3. Information Filtering

3.1 Where filtering occurs

The filtering performed by BARS will ensure that only the most relevant information is displayed to the user at a particular time. As illustrated in Figure 1, the filter regulates the flow of information from the Object to the User database. The filter regularly checks the database and, from the current user state (position and intent), it selects a set of objects which are passed to the User Database.

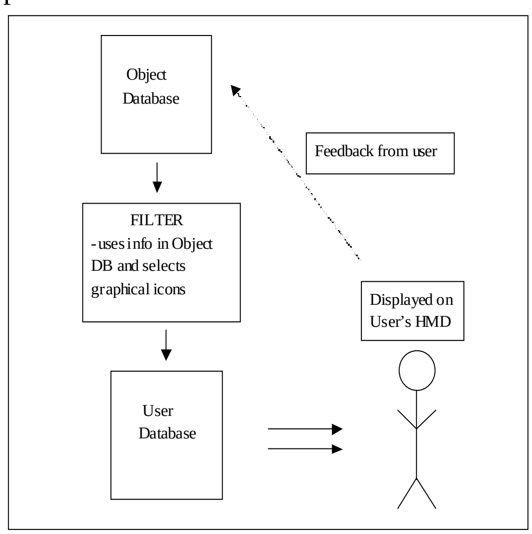

Figure 1: The Intelligent Filter fits in-between the Object Database and the User Database.

3.2 Displayed to user on the Helmet

The aim of the intelligent filter is to provide the graphics system with a list objects of that it should show the user upon his HMD. The first step in this process is to determine what the user can 'see'. Then, from this, determine what is relevant to the user based on their goal or mission. In the BARS system, there are six possible goals:

- FULL-ATTACK: Full attack on a particular target, such as a building or an enemy military installation.
- STRATEGIC-ATTACK: Strategic movement of troops towards a better strategic position in the urban environment. The attack is not centralized to a single enemy asset but looks for a better strategic situation for future operations.
- ROUTE: The mission of the user is to go from point A to point B in the urban environment in the shortest time possible. In a friendly environment such a route would be the most direct path, but in a combat scenario, it takes into account the position of enemy units, and zones where an enemy ambush is likely to take place. The route needs to be updated dynamically in real-time to reflect the user's current position and when new information on snipers is received, for example.
- STEALTH: The mission is to perform clandestine operations behind the enemy lines. These type of missions include intelligence gathering, sabotague, hostage rescue and indirect assault on enemy assets.
- RECON: The mission is to perform reconnaissance on a determined area of the urban environment. In this goal, the reconnaissance can be done on important military targets, such as military installations, arms caches, targets of tactical advantage and etc. Of importance for this mission are installations where civilians may have gathered, food stores and water deposits.
- FULL-RECON: The mission is to perform a more deep reconnaissance on a determined area of the urban environment. Besides concentrating on military and tactical targets done in RECON, FULL-RECON also takes into account civilian installa-

tions, and almost any other kind of building.

The filter employs a "two pass" mechanism to decide what information should be shown.

- The first mechanism is a physical one that only shows information that is 'close' to the user's current location.
- The second mechanism is logical or 'intelligent' one that only shows information to the user, which is 'relevant' to them.

3.3 Physical constraint

3.3.1 User's area of interest

The first step in this physical constraint is to determine the user's physical area of interest. This is the area that the user focuses upon. The user's current position is known at all times through the GPS/GLONASS[i] tracking system. The range that the user can see is determined by their goal. The following criterion is used, based upon the current goal:

- ROUTE – range is straight-line distance to the destination,
- RECON – medium level
- FULL-RECON – medium level,
- STEALTH – medium level
- STRATEGIC-ATTACK – low level
- FULL-ATTACK – low level.

This range can be overridden by the user at any time.

Given a range, a 3D volume of area can be determined centered around the user's position. For convenience a cube is used.

In the real-world, visual information is inhibited by real-world limitations such as walls. However, for the user in BARS, this is not a problem; BARS can provide "X-Ray" vision of the urban environment (based on available information in the database). This clearly highlights an additional advantage of the BARS system. Physical barriers do not limit information shown to the user

As a result of these considerations, an area of interest of the user is obtained.

3.3.2 Object's region of Influence (RI)

Definition

Every object in the environment has two cubes defining it:

- one enclosing the object's actual physical or geometric dimensions and
- one known as the Region of Influence.

To illustrate the difference between the two, consider a sniper. The geometric box would encompass the actual sniper; this would be extremely useful for attacking this sniper.

Based on database information, we can create a second box, known as the Region of Influence, centered at the sniper's position. In the case of the sniper, the range of this box is determined by the lethality range of the weapon (plus a small buffer) that the sniper is carrying. Even though the lethality range of the sniper's weapon would be the same, the box surrounding the sniper needs to move with the change in the sniper's position.

In order to provide this information to BARS, we have developed heuristics based on the object's type and importance to define a region of influence around an object. The region of influence actually defines the geographical zone where an object has an influence upon.

Object Type

The following criteria was used in order to calculate the Region of Influence (RI) of an object based upon its type:

- If an object contains no forces (ie moving objects), then its region-of-influence is equivalent to geometric region
- Otherwise if object is a force and is
  - Slow moving (eg sniper or ground troop), then its RI = lethality range of weapon plus a small percentage (eg 10%) or
  - Fast moving (eg tank or APC troop), then its RI = lethality range of weapon plus a large percentage (eg 20%).

Importance parameter

The second aspect to consider when defining a region of influence around an object, is the importance parameter associated with each object. In certain mission or goal, some objects will be important, while others will not be. For instance, in FULL-ATTACK attack there is no need to display the location of food

[i] GLONASS is the Russian equivalent of the US (NAVSTAR) GPS system.

supplies. In order to incorporate this information, each object has an importance parameter associated with it.

The importance parameter endeavours to capture the tactical knowledge of the objects in relation to its type, position and the goal of the user. For instance, a tall building provides a strategic advantage to the side that owns it, so it would be a high priority under the Strategic Attack mode; this is because it is a well known fact that it is very difficult for any kind of defensive manouveuer against an enemy in an elevated position. Thus in BARS we reflect this information by ensuring that the region of influence around the tall building is large as compared to that for a shorter building (for the Strategic Attack mode).

In the present prototype, the importance parameter is a binary vector which attempts to represent the importance on an object by asking some of the following questions:

- does the objects contain military assets?
- is the object or the zone of defensive tactical disadvantage, meaning that is a good place for an enemy ambush?
- is the object of offensive tactical advantage, such as a tall building ?

Using the Importance parameter, the Region of Influence of a particular object may be increased.

Calculating an Object's RI

Using the above two sets of criteria, every object in the environment will have a Region of Influence. This RI is centered at the object's physical location. The size of the region of influence is determined using the object's type and importance vector, as well as the user's current goal. In addition, for moving forces, the RI moves with the object. Thus, these regions in particular, need to be calculated in real-time in order to warn the user of possible danger. Figure 2 shows typical RI for some objects.

3.3.3 List of objects that the user can 'see'

Given the user's area of interest and the region-of-influence of every object in the environment, a search can now be made. If an object's region-of-influence intersects with the user's area of interest, then that object may be relevant to the user; see Figure 3. The result of this search is that a list of objects that intersect the user's area of interest is found.

However, this search has only considered the physical aspects of the user and the objects. This is because this search is based upon the user's area of interest and an object's region of influence. The next step involves determining whether the objects on this list should be displayed to the user.

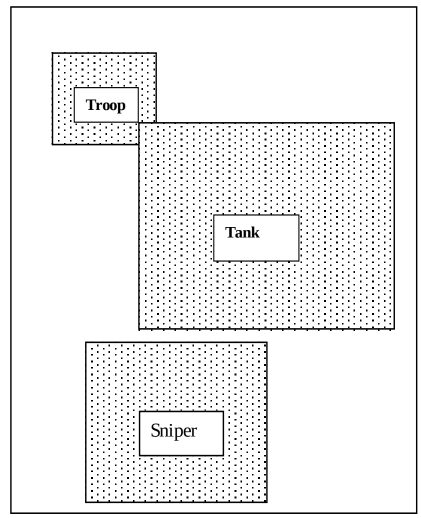

Figure 2:Regions of influence around particular objects.

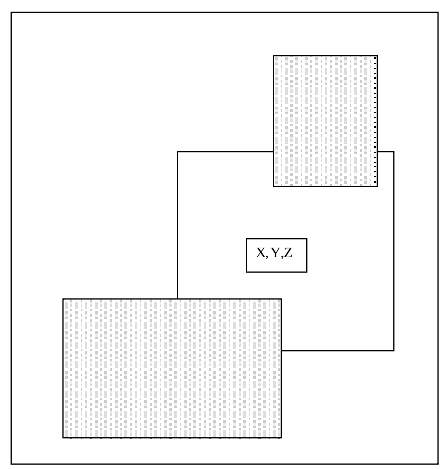

Figure 3: The user's area of interest intersecting some objects' region of influence.

3.4 Logical constraint

3.4.1 Strategy

The above technique of the physical constraints resulted in a list of potentially relevant objects being determined. This section describes a further refinement to this list so that only relevant objects are kept and displayed to the user. There are certain objects that need to be displayed to the user, regardless of the goal or their importance. These are discussed in the next section. Objects selected according to the current goal are discussed next.

3.4.2 High priority objects

There are objects that are always displayed on the user's HMD, regardless of the goal. These items are objects that are enemy controlled.

3.4.3 Objects relevant to goals.

In order to determine which objects to display to the user, the importance parameter is once again used to ensure that certain objects are displayed in certain modes. For instance, in an FULL-ATTACK mode, all friendly controlled objects are shown.

Besides determining which objects to show, this criterion also reflects which objects not to show. Using this closed-world assumption, the resulting list is tailored to the user's current situation. As illustrated in Figure 4, even though two objects were found to intersect the user's area of interest, only one was kept.

3.5 Overall technique

In order to determine the lists of objects to show a user at a particular time, both physical and logical constraints are used. The main steps for determining which objects should be shown to the user are listed below, in Figure 5. This algorithm lists the main steps used by the filter in order to determine which objects it should display to the user.

3.6 Timing of filter

The filter is the mechanism used in order to determine a list of objects or single object that should be shown to the user in the BARS prototype. The question now to ask is when should the filter be called.

The timing of the filter is based on an event-based loop. A full re-assessment of the current situation is performed when

- a new Goal is entered into the system or
- the user's region of interest changes (e.g., because the user has moved a prespecified distance *d* since the last filtering step) or
- a timeout event has occurred (the time since the last update exceeds some specified threshold *t*).

A more focussed filter is called when an object in the environment is added, deleted or modified in the Object Database. In this situation, only the effect of one particular object is gauged against the current objects shown to the user.

4. BARS Prototype

The BARS's hardware is illustrated in Figure 6 and is composed of the following off-the-shelf components:

- an Ashtech GG24-Surveyor (GPS receiver for position- only tracking),
- an InterSense IS300Pro (for orientation-only tracking),
- a Sony Glasstron Head-Mounted Display (HMD),
- a Dell Inspiron 7000 Notebook computer (main CPU and graphics engine) and
- a FreeWave Radio Modem (currently used just to broadcast GPS differential corrections).

The software is implemented using Java JDK 1.2 (for high level object management) and C (for high performance graphics rendering).

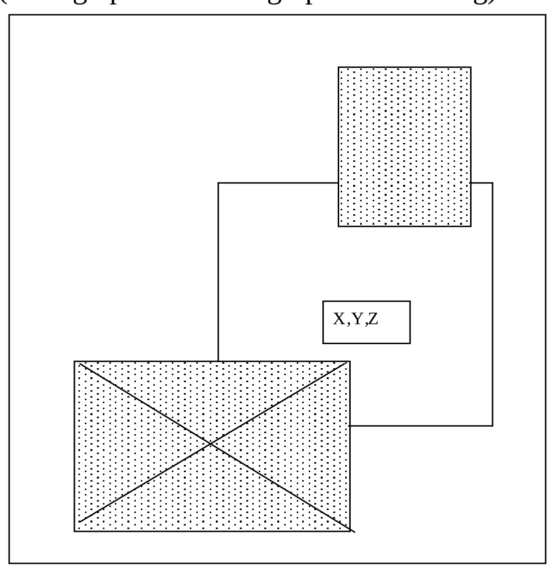

Figure 4: Eliminating an object from the list whose importance is not relevant in the current mode.

==

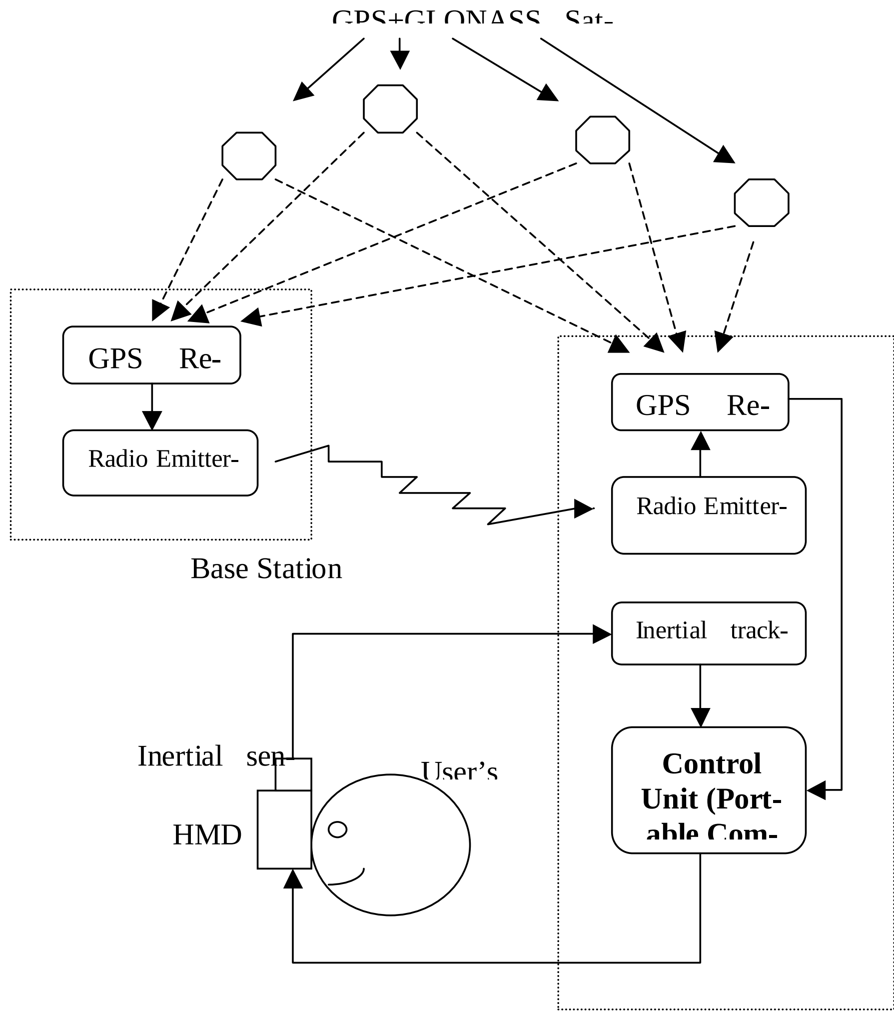

Calculate the *region of interest for the user*.
Calculate the *Region-of-influence (RI) for all the objects in the environment*

- calculate base RI using *object type*
- extend RI using *Importance parameter*

Determine objects whose *RI intersect User's area of interest*
Select objects that are relevant to the goal using *logical constraints*

- select all *enemy controlled objects*
- select all objects *pertinet* to *goal* using the *Importance Parameter* plus other considerations; keep or remove objects as necessary.

➔ LIST OF OBJECTS TO BE SHOWN TO USER

Figure 5: Main steps used by the filter.

5. Current Success and Future work

Figure 6:

The Prototype BARS Sys

l in its early stages. ng the idea behind the Importance Parameter and the structure of the Object Database. In addition, adding realistic military knowledge to the possible missions and the activities during those missions would enhance the realism of the system.

The filter itself represents the first step in an effort to develop an intelligent and autonomous Graphical Information Management System (GIMS) which will provide information in a hands-off manner to the user. The full GIMS capability extends the filter's concept of what should and should not be displayed by adding the extra dimension of determing *how* an object should be displayed. For example, "high priority" objects might require a different presentation style (e.g., through a different colour and/or the use of a supplemental figure audible cue such as an alarm).

6. Conclusions

This paper has described the Filtering mechnanism used by the Battlefield Augmented Reality System (BARS). The filter system dictates what should be shown to the user and when. It thus reduces the amount of possible information that the user in the field could possibly see, superimposed upon the real-world. In addition, the filter focuses the attention for the user. The filter can be thought of, as a data management system which takes in data at one end and send only relevant information through to the other end.

REFERENCES

[CFMOUT-97]

Concepts Division, Marine Corps Combat Development Command, "A Concept for Future
Military Operations on Urbanized Terrain," approved July 1997

[Gumm-98]

M.M.Gumm, W.P. Marshakm T.A. Branscome, M.Mc Wesler, D.J. Patton and L.L. Mullins, "A comparison of Solider Performance using current Land Navigation Equipment with Information integrated on a Helmet-Mounted Display, ARL Report ARL-TR-1604, DTIC Report 19980527081, April 1998.

[Julier-99]

S. Julier, S. Feiner and L. Rosenblum, "Augmented Reality as an Example of a Demanding Human-Centered System", First EC/NSF Advanced Research Workshop, 1-4 June 1999, France.